

2006: Celebrating 75 years of AI - History and Outlook: the Next 25 Years*

Jürgen Schmidhuber

TU Munich, Boltzmannstr. 3, 85748 Garching bei München, Germany &
IDSIA, Galleria 2, 6928 Manno (Lugano), Switzerland
juergen@idsia.ch - <http://www.idsia.ch/~juergen>

Abstract

When Kurt Gödel laid the foundations of theoretical computer science in 1931, he also introduced essential concepts of the theory of Artificial Intelligence (AI). Although much of subsequent AI research has focused on heuristics, which still play a major role in many practical AI applications, in the new millennium AI theory has finally become a full-fledged formal science, with important optimality results for embodied agents living in unknown environments, obtained through a combination of theory *à la* Gödel and probability theory. Here we look back at important milestones of AI history, mention essential recent results, and speculate about what we may expect from the next 25 years, emphasizing the significance of the ongoing dramatic hardware speedups, and discussing Gödel-inspired, self-referential, self-improving universal problem solvers.

1 Highlights of AI History—From Gödel to 2006

Gödel and Lilienfeld. In 1931, 75 years ago and just a few years after Julius Lilienfeld patented the transistor, Kurt Gödel laid the foundations of theoretical computer science (CS) with his work on universal formal languages and the limits of proof and computation [5]. He constructed formal systems allowing for self-referential statements that talk about themselves, in particular, about whether they can be derived from a set of given axioms through a computational theorem proving procedure. Gödel went on to construct statements that claim their own unprovability, to demonstrate that traditional math is either flawed in a certain algorithmic sense or contains unprovable but true statements.

Gödel's incompleteness result is widely regarded as the most remarkable achievement of 20th century mathematics, although some mathematicians say it is logic, not math, and others call it the fundamental result of theoretical computer science, a discipline that did not yet officially exist back then but was effectively created through

*Invited contribution to the Proceedings of the "50th Anniversary Summit of Artificial Intelligence" at Monte Verita, Ascona, Switzerland, 9-14 July 2006 (variant accepted for Springer's LNAI series)

Gödel's work. It had enormous impact not only on computer science but also on philosophy and other fields. In particular, since humans can “see” the truth of Gödel's unprovable statements, some researchers mistakenly thought that his results show that machines and Artificial Intelligences (AIs) will always be inferior to humans. Given the tremendous impact of Gödel's results on AI theory, it does make sense to date AI's beginnings back to his 1931 publication 75 years ago.

Zuse and Turing. In 1936 Alan Turing [37] introduced the *Turing machine* to reformulate Gödel's results and Alonzo Church's extensions thereof. TMs are often more convenient than Gödel's integer-based formal systems, and later became a central tool of CS theory. Simultaneously Konrad Zuse built the first working program-controlled computers (1935-1941), using the binary arithmetic and the *bits* of Gottfried Wilhelm von Leibniz (1701) instead of the more cumbersome decimal system used by Charles Babbage, who pioneered the concept of program-controlled computers in the 1840s, and tried to build one, although without success. By 1941, all the main ingredients of ‘modern’ computer science were in place, a decade after Gödel's paper, a century after Babbage, and roughly three centuries after Wilhelm Schickard, who started the history of automatic computing hardware by constructing the first non-program-controlled computer in 1623.

In the 1940s Zuse went on to devise the first high-level programming language (Plankalkül), which he used to write the first chess program. Back then chess-playing was considered an intelligent activity, hence one might call this chess program the first design of an AI program, although Zuse did not really implement it back then. Soon afterwards, in 1948, Claude Shannon [33] published information theory, recycling several older ideas such as Ludwig Boltzmann's entropy from 19th century statistical mechanics, and the *bit of information* (Leibniz, 1701).

Relays, Tubes, Transistors. Alternative instances of transistors, the concept pioneered and patented by Julius Edgar Lilienfeld (1920s) and Oskar Heil (1935), were built by William Shockley, Walter H. Brattain & John Bardeen (1948: point contact transistor) as well as Herbert F. Mataré & Heinrich Walker (1948, exploiting transconductance effects of germanium diodes observed in the *Luftwaffe* during WW-II). Today most transistors are of the field-effect type *à la* Lilienfeld & Heil. In principle a switch remains a switch no matter whether it is implemented as a relay or a tube or a transistor, but transistors switch faster than relays (Zuse, 1941) and tubes (Colossus, 1943; ENIAC, 1946). This eventually led to significant speedups of computer hardware, which was essential for many subsequent AI applications.

The I in AI. In 1950, some 56 years ago, Turing invented a famous subjective test to decide whether a machine or something else is intelligent. 6 years later, and 25 years after Gödel's paper, John McCarthy finally coined the term “AI”. 50 years later, in 2006, this prompted some to celebrate the 50th birthday of AI, but this chapter's title should make clear that its author cannot agree with this view—it is the thing that counts, not its name.

Roots of Probability-Based AI. In the 1960s and 1970s Ray Solomonoff combined theoretical CS and probability theory to establish a general theory of universal inductive inference and predictive AI [35] closely related to the concept of Kolmogorov complexity [14]. His theoretically optimal predictors and their Bayesian learning algorithms only assume that the observable reactions of the environment in response to cer-

tain action sequences are sampled from an unknown probability distribution contained in a set M of all enumerable distributions. That is, given an observation sequence we only assume there exists a computer program that can compute the probabilities of the next possible observations. This includes all scientific theories of physics, of course. Since we typically do not know this program, we predict using a weighted sum ξ of *all* distributions in M , where the sum of the weights does not exceed 1. It turns out that this is indeed the best one can possibly do, in a very general sense [11, 35]. Although the universal approach is practically infeasible since M contains infinitely many distributions, it does represent the first sound and general theory of optimal prediction based on experience, identifying the limits of both human and artificial predictors, and providing a yardstick for all prediction machines to come.

AI vs Astrology? Unfortunately, failed prophecies of human-level AI with just a tiny fraction of the brain’s computing power discredited some of the AI research in the 1960s and 70s. Many theoretical computer scientists actually regarded much of the field with contempt for its perceived lack of hard theoretical results. ETH Zurich’s Turing award winner and creator of the PASCAL programming language, Niklaus Wirth, did not hesitate to link AI to astrology. Practical AI of that era was dominated by rule-based expert systems and Logic Programming. That is, despite Solomonoff’s fundamental results, a main focus of that time was on logical, deterministic deduction of facts from previously known facts, as opposed to (probabilistic) induction of hypotheses from experience.

Evolution, Neurons, Ants. Largely unnoticed by mainstream AI gurus of that era, a biology-inspired type of AI emerged in the 1960s when Ingo Rechenberg pioneered the method of artificial evolution to solve complex optimization tasks [22], such as the design of optimal airplane wings or combustion chambers of rocket nozzles. Such methods (and later variants thereof, e.g., Holland [10] (1970s), often gave better results than classical approaches. In the following decades, other types of “sub-symbolic” AI also became popular, especially neural networks. Early neural net papers include those of McCulloch & Pitts, 1940s (linking certain simple neural nets to old and well-known, simple mathematical concepts such as linear regression); Minsky & Papert [17] (temporarily discouraging neural network research), Kohonen [12], Amari, 1960s; Werbos [40], 1970s; and many others in the 1980s. Orthogonal approaches included fuzzy logic (Zadeh, 1960s), Rissanen’s practical variants [23] of Solomonoff’s universal method, “representation-free” AI (Brooks [2]), Artificial Ants (Dorigo & Gambardella [4], 1990s), statistical learning theory (in less general settings than those studied by Solomonoff) & support vector machines (Vapnik [38] and others). As of 2006, this alternative type of AI research is receiving more attention than “Good Old-Fashioned AI” (GOFAI).

Mainstream AI Marries Statistics. A dominant theme of the 1980s and 90s was the marriage of mainstream AI and old concepts from probability theory. Bayes networks, Hidden Markov Models, and numerous other probabilistic models found wide applications ranging from pattern recognition, medical diagnosis, data mining, machine translation, robotics, etc.

Hardware Outshining Software: Humanoids, Robot Cars, Etc. In the 1990s and 2000s, much of the progress in practical AI was due to better hardware, getting roughly 1000 times faster per Euro per decade. In 1995, a fast vision-based robot car

by Ernst Dickmanns (whose team built the world's first reliable robot cars in the early 1980s with the help of Mercedes-Benz, e. g., [3]) autonomously drove 1000 miles from Munich to Denmark and back, in traffic at up to 120 mph, automatically passing other cars (a safety driver took over only rarely in critical situations). Japanese labs (Honda, Sony) and Pfeiffer's lab at TU Munich built famous humanoid walking robots. Engineering problems often seemed more challenging than AI-related problems.

Another source of progress was the dramatically improved access to all kinds of data through the WWW, created by Tim Berners-Lee at the European particle collider CERN (Switzerland) in 1990. This greatly facilitated and encouraged all kinds of "intelligent" data mining applications. However, there were few if any obvious fundamental algorithmic breakthroughs; improvements / extensions of already existing algorithms seemed less impressive and less crucial than hardware advances. For example, chess world champion Kasparov was beaten by a fast IBM computer running a fairly standard algorithm. Rather simple but computationally expensive probabilistic methods for speech recognition, statistical machine translation, computer vision, optimization, virtual realities etc. started to become feasible on PCs, mainly because PCs had become 1000 times more powerful within a decade or so.

2006. As noted by Stefan Artmann (personal communication, 2006), today's AI textbooks seem substantially more complex and less unified than those of several decades ago, e. g., [18], since they have to cover so many apparently quite different subjects. There seems to be a need for a new unifying view of intelligence. In the author's opinion this view already exists, as will be discussed below.

2 Subjective Selected Highlights of Present AI

The more recent some event, the harder it is to judge its long-term significance. But this biased author thinks that the most important thing that happened recently in AI is the begin of a transition from a heuristics-dominated science (e.g., [24]) to a real formal science. Let us elaborate on this topic.

2.1 The Two Ways of Making a Dent in AI Research

There are at least two convincing ways of doing AI research: **(1)** construct a (possibly heuristic) machine or algorithm that somehow (it does not really matter how) solves a previously unsolved interesting problem, such as beating the best human player of *Go* (success will outshine any lack of theory). Or **(2)** prove that a particular novel algorithm is optimal for an important class of AI problems.

It is the nature of heuristics (case **(1)**) that they lack staying power, as they may soon get replaced by next year's even better heuristics. Theorems (case **(2)**), however, are for eternity. That's why formal sciences prefer theorems.

For example, probability theory became a formal science centuries ago, and totally formal in 1933 with Kolmogorov's axioms [13], shortly after Gödel's paper [5]. Old but provably optimal techniques of probability theory are still in every day's use, and in fact highly significant for modern AI, while many initially successful heuristic approaches eventually became unfashionable, of interest mainly to the historians of the field.

2.2 No Brain Without a Body / AI Becoming a Formal Science

Heuristic approaches will continue to play an important role in many AI applications, to the extent they empirically outperform competing methods. But like with all young sciences at the transition point between an early intuition-dominated and a later formal era, the importance of mathematical optimality theorems is growing quickly. Progress in the formal era, however, is and will be driven by a different breed of researchers, a fact that is not necessarily universally enjoyed and welcomed by all the earlier pioneers.

Today the importance of embodied, embedded AI is almost universally acknowledged (e. g., [20]), as obvious from frequently overheard remarks such as “let the physics compute” and “no brain without a body.” Many present AI researchers focus on real robots living in real physical environments. To some of them the title of this subsection may seem oxymoronic: the extension of AI into the realm of the physical body seems to be a step away from formalism. But the new millennium’s formal point of view is actually taking this step into account in a very general way, through the first mathematical theory of universal embedded AI, combining “old” theoretical computer science and “ancient” probability theory to derive optimal behavior for embedded, embodied rational agents living in unknown but learnable environments. More on this below.

2.3 What’s the I in AI? What is Life? Etc.

Before we proceed, let us clarify what we are talking about. Shouldn’t researchers on Artificial Intelligence (AI) and Artificial Life (AL) agree on basic questions such as: What is Intelligence? What is Life? Interestingly they don’t.

Are Cars Alive? For example, AL researchers often offer definitions of life such as: it must reproduce, evolve, etc. Cars are alive, too, according to most of these definitions. For example, cars evolve and multiply. They need complex environments with car factories to do so, but living animals also need complex environments full of chemicals and other animals to reproduce — the DNA information by itself does not suffice. There is no obvious fundamental difference between an organism whose self-replication information is stored in its DNA, and a car whose self-replication information is stored in a car builder’s manual in the glove compartment. To copy itself, the organism needs its mother’s womb plus numerous other objects and living beings in its environment (such as trillions of bacteria inside and outside of the mother’s body). The car needs iron mines and car part factories and human workers.

What is Intelligence? If we cannot agree on what’s life, or, for that matter, love, or consciousness (another fashionable topic), how can there be any hope to define intelligence? Turing’s definition (1950, 19 years after Gödel’s paper) was totally subjective: intelligent is what convinces me that it is intelligent while I am interacting with it. Fortunately, however, there are more formal and less subjective definitions.

2.4 Formal AI Definitions

Popper said: all life is problem solving [21]. Instead of defining intelligence in Turing’s rather vague and subjective way we define intelligence with respect to the abilities of

universal optimal problem solvers.

Consider a learning robotic agent with a single life which consists of discrete cycles or time steps $t = 1, 2, \dots, T$. Its total lifetime T may or may not be known in advance. In what follows, the value of any time-varying variable Q at time t ($1 \leq t \leq T$) will be denoted by $Q(t)$, the ordered sequence of values $Q(1), \dots, Q(t)$ by $Q(\leq t)$, and the (possibly empty) sequence $Q(1), \dots, Q(t-1)$ by $Q(< t)$.

At any given t the robot receives a real-valued input vector $x(t)$ from the environment and executes a real-valued action $y(t)$ which may affect future inputs; at times $t < T$ its goal is to maximize future success or *utility*

$$u(t) = E_\mu \left[\sum_{\tau=t+1}^T r(\tau) \mid h(\leq t) \right], \quad (1)$$

where $r(t)$ is an additional real-valued reward input at time t , $h(t)$ the ordered triple $[x(t), y(t), r(t)]$ (hence $h(\leq t)$ is the known history up to t), and $E_\mu(\cdot \mid \cdot)$ denotes the conditional expectation operator with respect to some possibly unknown distribution μ from a set M of possible distributions. Here M reflects whatever is known about the possibly probabilistic reactions of the environment. For example, M may contain all computable distributions [11, 35]. Note that unlike in most previous work by others [36], there is just one life, no need for predefined repeatable trials, no restriction to Markovian interfaces between sensors and environment, and the utility function implicitly takes into account the expected remaining lifespan $E_\mu(T \mid h(\leq t))$ and thus the possibility to extend it through appropriate actions [29].

Any formal problem or sequence of problems can be encoded in the reward function. For example, the reward functions of many living or robotic beings cause occasional hunger or pain or pleasure signals etc. At time t an optimal AI will make the best possible use of experience $h(\leq t)$ to maximize $u(t)$. But how?

2.5 Universal, Mathematically Optimal, But Incomputable AI

Unbeknownst to many traditional AI researchers, there is indeed an extremely general “best” way of exploiting previous experience. At any time t , the recent theoretically optimal yet practically infeasible reinforcement learning (RL) algorithm AIXI [11] uses Solomonoff’s above-mentioned universal prediction scheme to select those action sequences that promise maximal future reward up to some horizon, given the current data $h(\leq t)$. Using a variant of Solomonoff’s universal probability mixture ξ , in cycle $t+1$, AIXI selects as its next action the first action of an action sequence maximizing ξ -predicted reward up to the horizon. Hutter’s recent work [11] demonstrated AIXI’s optimal use of observations as follows. The Bayes-optimal policy p^ξ based on the mixture ξ is self-optimizing in the sense that its average utility value converges asymptotically for all $\mu \in \mathcal{M}$ to the optimal value achieved by the (infeasible) Bayes-optimal policy p^μ which knows μ in advance. The necessary condition that \mathcal{M} admits self-optimizing policies is also sufficient.

Of course one cannot claim the old AI is devoid of formal research! The recent approach above, however, goes far beyond previous formally justified but very limited AI-related approaches ranging from linear perceptrons [17] to the A^* -algorithm [18].

It provides, for the first time, a mathematically sound theory of general AI and optimal decision making based on experience, identifying the limits of both human and artificial intelligence, and a yardstick for any future, scaled-down, practically feasible approach to general AI.

2.6 Optimal Curiosity and Creativity

No theory of AI will be convincing if it does not explain curiosity and creativity, which many consider as important ingredients of intelligence. We can provide an explanation in the framework of optimal reward maximizers such as those from the previous subsection.

It is possible to come up with theoretically optimal ways of improving the predictive world model of a curious robotic agent [28], extending earlier ideas on how to implement artificial curiosity [25]: *The rewards of an optimal reinforcement learner are the predictor's improvements* on the observation history so far. They encourage the reinforcement learner to produce action sequences that cause the creation and the learning of new, previously unknown regularities in the sensory input stream. It turns out that art and creativity can be explained as by-products of such intrinsic curiosity rewards: good observer-dependent art deepens the observer's insights about this world or possible worlds, connecting previously disconnected patterns in an initially surprising way that eventually becomes known and boring. While previous attempts at describing what is satisfactory art or music were informal, this work permits the first *technical, formal* approach to understanding the nature of art and creativity [28].

2.7 Computable, Asymptotically Optimal General Problem Solver

Using the Speed Prior [26] one can scale down the universal approach above such that it becomes computable. In what follows we will mention general methods whose optimality criteria explicitly take into account the computational costs of prediction and decision making—compare [15].

The recent asymptotically optimal search algorithm for *all* well-defined problems [11] allocates part of the total search time to searching the space of proofs for provably correct candidate programs with provable upper runtime bounds; at any given time it focuses resources on those programs with the currently best proven time bounds. The method is as fast as the initially unknown fastest problem solver for the given problem class, save for a constant slowdown factor of at most $1 + \epsilon$, $\epsilon > 0$, and an additive constant that does not depend on the problem instance!

Is this algorithm then the *holy grail* of computer science? Unfortunately not quite, since the additive constant (which disappears in the $O()$ -notation of theoretical CS) may be huge, and practical applications may not ignore it. This motivates the next section, which addresses all kinds of formal optimality (not just asymptotic optimality).

2.8 Fully Self-Referential, Self-Improving Gödel Machine

We may use Gödel's self-reference trick to build a universal general, fully self-referential, self-improving, optimally efficient problem solver [29]. A Gödel Machine is a com-

puter whose original software includes axioms describing the hardware and the original software (this is possible without circularity) plus whatever is known about the (probabilistic) environment plus some formal goal in form of an arbitrary user-defined utility function, e.g., cumulative future expected reward in a sequence of optimization tasks - see equation (1). The original software also includes a proof searcher which uses the axioms (and possibly an online variant of Levin’s universal search [15]) to systematically make pairs (“proof”, “program”) until it finds a proof that a rewrite of the original software through “program” will increase utility. The machine can be designed such that each self-rewrite is necessarily globally optimal in the sense of the utility function, even those rewrites that destroy the proof searcher [29].

2.9 Practical Algorithms for Program Learning

The theoretically optimal universal methods above are optimal in ways that do not (yet) immediately yield practically feasible general problem solvers, due to possibly large initial overhead costs. Which are today’s practically most promising extensions of traditional machine learning?

Since virtually all realistic sensory inputs of robots and other cognitive systems are sequential by nature, the future of machine learning and AI in general depends on progress in in sequence processing as opposed to the traditional processing of stationary input patterns. To narrow the gap between learning abilities of humans and machines, we will have to study how to learn general algorithms instead of such reactive mappings. Most traditional methods for learning time series and mappings from sequences to sequences, however, are based on simple time windows: one of the numerous feedforward ML techniques such as feedforward neural nets (NN) [1] or support vector machines [38] is used to map a restricted, fixed time window of sequential input values to desired target values. Of course such approaches are bound to fail if there are temporal dependencies exceeding the time window size. Large time windows, on the other hand, yield unacceptable numbers of free parameters.

Presently studied, rather general sequence learners include certain probabilistic approaches and especially recurrent neural networks (RNNs), e.g., [19]. RNNs have adaptive feedback connections that allow them to learn mappings from input sequences to output sequences. They can implement any sequential, algorithmic behavior implementable on a personal computer. In gradient-based RNNs, however, we can *differentiate our wishes with respect to programs*, to obtain a search direction in algorithm space. RNNs are biologically more plausible and computationally more powerful than other adaptive models such as Hidden Markov Models (HMMs - no continuous internal states), feedforward networks & Support Vector Machines (no internal states at all). For several reasons, however, the first RNNs could not learn to look far back into the past. This problem was overcome by RNNs of the *Long Short-Term Memory* type (LSTM), currently the most powerful and practical supervised RNN architecture for many applications, trainable either by gradient descent [9] or evolutionary methods [32], occasionally profiting from a marriage with probabilistic approaches [8].

Unsupervised RNNs that learn without a teacher to control physical processes or robots frequently use evolutionary algorithms [10, 22] to learn appropriate programs (RNN weight matrices) through trial and error [41]. Recent work brought progress

through a focus on reducing search spaces by co-evolving the comparatively small weight vectors of individual recurrent neurons [7]. Such RNNs can learn to create memories of important events, solving numerous RL / optimization tasks unsolvable by traditional RL methods [6, 7]. They are among the most promising methods for practical program learning, and currently being applied to the control of sophisticated robots such as the walking biped of TU Munich [16].

3 The Next 25 Years

Where will AI research stand in 2031, 25 years from now, 100 years after Gödel’s ground-breaking paper [5], some 200 years after Babbage’s first designs, some 400 years after the first automatic calculator by Schickard (and some 2000 years after the crucifixion of the man whose birth year anchors the Western calendar)?

Trivial predictions are those that just naively extrapolate the current trends, such as: computers will continue to get faster by a factor of roughly 1000 per decade; hence they will be at least a million times faster by 2031. According to frequent estimates, current supercomputers achieve roughly 1 percent of the raw computational power of a human brain, hence those of 2031 will have 10,000 “brain powers”; and even cheap devices will achieve many brain powers. Many tasks that are hard for today’s software on present machines will become easy without even fundamentally changing the algorithms. This includes numerous pattern recognition and control tasks arising in factories of many industries, currently still employing humans instead of robots.

Will theoretical advances and practical software keep up with the hardware development? We are convinced they will. As discussed above, the new millennium has already brought fundamental new insights into the problem of constructing theoretically optimal rational agents or universal AIs, even if those do not yet immediately translate into practically feasible methods. On the other hand, on a more practical level, there has been rapid progress in learning algorithms for agents interacting with a dynamic environment, autonomously discovering true sequence-processing, problem-solving programs, as opposed to the reactive mappings from stationary inputs to outputs studied in most of traditional machine learning research. In the author’s opinion the above-mentioned theoretical and practical strands are going to converge. In conjunction with the ongoing hardware advances this will yield non-universal but nevertheless rather general artificial problem-solvers whose capabilities will exceed those of most if not all humans in many domains of commercial interest. This may seem like a bold prediction to some, but it is actually a trivial one as there are so many experts who would agree with it.

Nontrivial predictions are those that anticipate truly unexpected, revolutionary breakthroughs. By definition, these are hard to predict. For example, in 1985 only very few scientists and science fiction authors predicted the WWW revolution of the 1990s. The few who did were not influential enough to make a significant part of humanity believe in their predictions and prepare for their coming true. Similarly, after the latest stock market crash one can always find with high probability some “prophet in the desert” who predicted it in advance, but had few if any followers until the crash really occurred.

Truly nontrivial predictions are those that most will not believe until they come true. We will mostly restrict ourselves to trivial predictions like those above and refrain from too much speculation in form of nontrivial ones. However, we may have a look at previous unexpected scientific breakthroughs and try to discern a pattern, a pattern that may not allow us to precisely predict the details of the next revolution but at least its timing.

3.1 A Pattern in the History of Revolutions?

Let us put the AI-oriented developments [27] discussed above in a broader context, and look at the history of major scientific revolutions and essential historic developments (that is, the subjects of the major chapters in history books) since the beginnings of modern man over 40,000 years ago [30, 31]. Amazingly, they seem to match a binary logarithmic scale marking exponentially declining temporal intervals [31], each half the size of the previous one, and measurable in terms of powers of 2 multiplied by a human lifetime (roughly 80 years—throughout recorded history many individuals have reached this age, although the average lifetime often was shorter, mostly due to high children mortality). It looks as if history itself will *converge* in a historic singularity or Omega point Ω around 2040 (the term *historic singularity* is apparently due to Stanislaw Ulam (1950s) and was popularized by Vernor Vinge [39] in the 1990s). To convince yourself of history’s convergence, associate an error bar of not much more than 10 percent with each date below:

1. $\Omega - 2^9$ lifetimes: modern humans start colonizing the world from Africa
2. $\Omega - 2^8$ lifetimes: bow and arrow invented; hunting revolution
3. $\Omega - 2^7$ lifetimes: invention of agriculture; first permanent settlements; beginnings of civilization
4. $\Omega - 2^6$ lifetimes: first high civilizations (Sumeria, Egypt), and the most important invention of recorded history, namely, the one that made recorded history possible: writing
5. $\Omega - 2^5$ lifetimes: the ancient Greeks invent democracy and lay the foundations of Western science and art and philosophy, from algorithmic procedures and formal proofs to anatomically perfect sculptures, harmonic music, and organized sports. Old Testament written (basis of Judaism, Christianity, Islam); major Asian religions founded. High civilizations in China, origin of the first calculation tools, and India, origin of alphabets and the zero
6. $\Omega - 2^4$ lifetimes: bookprint (often called the most important invention of the past 2000 years) invented in China. Islamic science and culture start spreading across large parts of the known world (this has sometimes been called the most important event between Antiquity and the age of discoveries)
7. $\Omega - 2^3$ lifetimes: the Mongolian Empire, the largest and most dominant empire ever (possibly including most of humanity and the world economy), stretches

across Asia from Korea all the way to Germany. Chinese fleets and later also European vessels start exploring the world. Gun powder and guns invented in China. Renaissance and Western bookprint (often called the most influential invention of the past 1000 years) and subsequent Reformation in Europe. Begin of the Scientific Revolution

8. $\Omega - 2^2$ lifetimes: Age of enlightenment and rational thought in Europe. Massive progress in the sciences; first flying machines; first steam engines prepare the industrial revolution
9. $\Omega - 2$ lifetimes: Second industrial revolution based on combustion engines, cheap electricity, and modern chemistry. Birth of modern medicine through the germ theory of disease; genetic and evolution theory. European colonialism at its short-lived peak
10. $\Omega - 1$ lifetime: modern post-World War II society and pop culture emerges; superpower stalemate based on nuclear deterrence. The 20th century super-exponential population explosion (from 1.6 billion to 6 billion people, mainly due to the Haber-Bosch process [34]) is at its peak. First spacecraft and commercial computers; DNA structure unveiled
11. $\Omega - 1/2$ lifetime (now): for the first time in history most of the most destructive weapons are dismantled, after the Cold War's peaceful end. 3rd industrial revolution based on personal computers and the World Wide Web. A mathematical theory of universal AI emerges (see sections above) - will this be considered a milestone in the future?
12. $\Omega - 1/4$ lifetime: This point will be reached around 2020. By then many computers will have substantially more raw computing power than human brains.
13. $\Omega - 1/8$ lifetime (100 years after Gödel's paper): will practical variants of Gödel machines start a runaway evolution of continually self-improving superminds way beyond human imagination, causing far more unpredictable revolutions in the final decade before Ω than during all the millennia before?
14. ...

The following disclosure should help the reader to take this list with a grain of salt though. The author, who admits being very interested in witnessing Ω , was born in 1963, and therefore perhaps should not expect to live long past 2040. This may motivate him to uncover certain historic patterns that fit his desires, while ignoring other patterns that do not. Perhaps there even is a general rule for both the individual memory of single humans and the collective memory of entire societies and their history books: constant amounts of memory space get allocated to exponentially larger, adjacent time intervals further and further into the past. Maybe that's why there has never been a shortage of prophets predicting that the end is near - the important events according to one's own view of the past always seem to accelerate exponentially. See [31] for a more thorough discussion of this possibility.

References

- [1] C. M. Bishop. *Neural networks for pattern recognition*. Oxford University Press, 1995.
- [2] R. A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 569–595, 1991.
- [3] E. D. Dickmanns, R. Behringer, D. Dickmanns, T. Hildebrandt, M. Maurer, F. Thomanek, and J. Schiehlen. The seeing passenger car 'VaMoRs-P'. In *Proc. Int. Symp. on Intelligent Vehicles '94, Paris*, pages 68–73, 1994.
- [4] M. Dorigo, G. Di Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, 1999.
- [5] K. Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.
- [6] F. Gomez, J. Schmidhuber, and R. Miikkulainen. Efficient non-linear control through neuroevolution. In *ECML 2006: Proceedings of the 17th European Conference on Machine Learning*. Springer, 2006.
- [7] F. J. Gomez and R. Miikkulainen. Active guidance for a finless rocket using neuroevolution. In *Proc. GECCO 2003, Chicago*, 2003. *Winner of Best Paper Award in Real World Applications. Gomez is working at IDSIA on a CSEM grant to J. Schmidhuber.*
- [8] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural nets. In *ICML '06: Proceedings of the International Conference on Machine Learning*, 2006.
- [9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [10] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.
- [11] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2004. (On J. Schmidhuber's SNF grant 20-61847).
- [12] T. Kohonen. *Self-Organization and Associative Memory*. Springer, second edition, 1988.
- [13] A. N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933.
- [14] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–11, 1965.

- [15] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.
- [16] S. Lohmeier, K. Loeffler, M. Gienger, H. Ulbrich, and F. Pfeiffer. Sensor system and trajectory control of a biped robot. In *Proc. 8th IEEE International Workshop on Advanced Motion Control (AMC'04), Kawasaki, Japan*, pages 393–398, 2004.
- [17] M. Minsky and S. Papert. *Perceptrons*. Cambridge, MA: MIT Press, 1969.
- [18] N. J. Nilsson. *Principles of artificial intelligence*. Morgan Kaufmann, San Francisco, CA, USA, 1980.
- [19] B. A. Pearlmutter. Gradient calculations for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks*, 6(5):1212–1228, 1995.
- [20] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, 2001.
- [21] K. R. Popper. *All Life Is Problem Solving*. Routledge, London, 1999.
- [22] I. Rechenberg. Evolutionsstrategie - Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Dissertation, 1971. Published 1973 by Fromman-Holzboog.
- [23] J. Rissanen. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.
- [24] P. S. Rosenbloom, J. E. Laird, and A. Newell. *The SOAR Papers*. MIT Press, 1993.
- [25] J. Schmidhuber. Curious model-building control systems. In *Proceedings of the International Joint Conference on Neural Networks, Singapore*, volume 2, pages 1458–1463. IEEE press, 1991.
- [26] J. Schmidhuber. The Speed Prior: a new simplicity measure yielding near-optimal computable predictions. In J. Kivinen and R. H. Sloan, editors, *Proceedings of the 15th Annual Conference on Computational Learning Theory (COLT 2002)*, Lecture Notes in Artificial Intelligence, pages 216–228. Springer, Sydney, Australia, 2002.
- [27] J. Schmidhuber. Artificial Intelligence - history highlights and outlook: AI maturing and becoming a real formal science, 2006. <http://www.idsia.ch/~juergen/ai.html>.
- [28] J. Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187, 2006.
- [29] J. Schmidhuber. Gödel machines: fully self-referential optimal universal problem solvers. In B. Goertzel and C. Pennachin, editors, *Artificial General Intelligence*, pages 199–226. Springer Verlag, 2006.

- [30] J. Schmidhuber. Is history converging? Again?, 2006.
<http://www.idsia.ch/~juergen/history.html>.
- [31] J. Schmidhuber. New millennium AI and the convergence of history. In W. Duch and J. Mandziuk, editors, *Challenges to Computational Intelligence*. Springer, in press, 2006. Also available as TR IDSIA-04-03, cs.AI/0302012.
- [32] J. Schmidhuber, D. Wierstra, M. Gagliolo, and F. Gomez. Training recurrent networks by EVOLINO. *Neural Computation*, 19(3):757–779, 2007.
- [33] C. E. Shannon. A mathematical theory of communication (parts I and II). *Bell System Technical Journal*, XXVII:379–423, 1948.
- [34] V. Smil. Detonator of the population explosion. *Nature*, 400:415, 1999.
- [35] R. J. Solomonoff. Complexity-based induction systems. *IEEE Transactions on Information Theory*, IT-24(5):422–432, 1978.
- [36] R. Sutton and A. Barto. *Reinforcement learning: An introduction*. Cambridge, MA, MIT Press, 1998.
- [37] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, Series 2*, 41:230–267, 1936.
- [38] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer, New York, 1995.
- [39] V. Vinge. The coming technological singularity, 1993. VISION-21 Symposium sponsored by NASA Lewis Research Center, and Whole Earth Review, Winter issue.
- [40] P. J. Werbos. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard University, 1974.
- [41] Xin Yao. A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 4:203–222, 1993.